Internet Engineering Task Force (IETF)
Request for Comments: 7258
BCP: 188
Category: Best Current Practice
ISSN: 2070-1721
S. Farrell
Trinity College Dublin
H. Tschofenig
ARM Ltd.
May 2014

Pervasive Monitoring Is an Attack

Abstract

Pervasive monitoring is a technical attack that should be mitigated in the design of IETF protocols, where possible.

Status of This Memo

Copyright Notice

1. Pervasive Monitoring Is a Widespread Attack on Privacy

Pervasive Monitoring (PM) is widespread (and often covert) surveillance through intrusive gathering of protocol artefacts, including application content, or protocol metadata such as headers. Active or passive wiretaps and traffic analysis, (e.g., correlation, timing or measuring packet sizes), or subverting the cryptographic keys used to secure protocols can also be used as part of pervasive monitoring. PM is distinguished by being indiscriminate and very large scale, rather than by introducing new types of technical compromise.

The IETF community's technical assessment is that PM is an attack on the privacy of Internet users and organisations. The IETF community has expressed strong agreement that PM is an attack that needs to be mitigated where possible, via the design of protocols that make PM significantly more expensive or infeasible. Pervasive monitoring was discussed at the technical plenary of the November 2013 IETF meeting [IETF88Plenary] and then through extensive exchanges on IETF mailing lists. This document records the IETF community's consensus and establishes the technical nature of PM.

The term "attack" is used here in a technical sense that differs somewhat from common English usage. In common English usage, an attack is an aggressive action perpetrated by an opponent, intended to enforce the opponent's will on the attacked party. The term is used here to refer to behavior that subverts the intent of communicating parties without the agreement of those parties. An attack may change the content of the communication, record the content or external characteristics of the communication, or through correlation with other communication events, reveal information the parties did not intend to be revealed. It may also have other effects that similarly subvert the intent of a communicator. [RFC4949] contains a more complete definition for the term "attack". We also use the term in the singular here, even though PM in reality may consist of a multifaceted set of coordinated attacks.

In particular, the term "attack", used technically, implies nothing about the motivation of the actor mounting the attack. The motivation for PM can range from non-targeted nation-state surveillance, to legal but privacy-unfriendly purposes by commercial enterprises, to illegal actions by criminals. The same techniques to achieve PM can be used regardless of motivation. Thus, we cannot defend against the most nefarious actors while allowing monitoring by other actors no matter how benevolent some might consider them to be, since the actions required of the attacker are indistinguishable from other attacks. The motivation for PM is, therefore, not relevant for how PM is mitigated in IETF protocols.

2. The IETF Will Work to Mitigate Pervasive Monitoring

"Mitigation" is a technical term that does not imply an ability to completely prevent or thwart an attack. Protocols that mitigate PM will not prevent the attack but can significantly change the threat. (See the diagram on page 24 of RFC 4949 for how the terms "attack" and "threat" are related.) This can significantly increase the cost of attacking, force what was covert to be overt, or make the attack more likely to be detected, possibly later.

IETF standards already provide mechanisms to protect Internet communications and there are guidelines [RFC3552] for applying these in protocol design. But those standards generally do not address PM, the confidentiality of protocol metadata, countering traffic analysis, or data minimisation. In all cases, there will remain some privacy-relevant information that is inevitably disclosed by protocols. As technology advances, techniques that were once only available to extremely well-funded actors become more widely accessible. Mitigating PM is therefore a protection against a wide range of similar attacks.

It is therefore timely to revisit the security and privacy properties of our standards. The IETF will work to mitigate the technical aspects of PM, just as we do for protocol vulnerabilities in general. The ways in which IETF protocols mitigate PM will change over time as mitigation and attack techniques evolve and so are not described here.

Those developing IETF specifications need to be able to describe how they have considered PM, and, if the attack is relevant to the work to be published, be able to justify related design decisions. This does not mean a new "pervasive monitoring considerations" section is needed in IETF documentation. It means that, if asked, there needs to be a good answer to the question "Is pervasive monitoring relevant to this work and if so, how has it been considered?"

In particular, architectural decisions, including which existing technology is reused, may significantly impact the vulnerability of a protocol to PM. Those developing IETF specifications therefore need to consider mitigating PM when making architectural decisions. Getting adequate, early review of architectural decisions including whether appropriate mitigation of PM can be made is important. Revisiting these architectural decisions late in the process is very costly.

While PM is an attack, other forms of monitoring that might fit the definition of PM can be beneficial and not part of any attack, e.g., network management functions monitor packets or flows and anti-spam

mechanisms need to see mail message content. Some monitoring can even be part of the mitigation for PM, for example, certificate transparency [RFC6962] involves monitoring Public Key Infrastructure in ways that could detect some PM attack techniques. However, there is clear potential for monitoring mechanisms to be abused for PM, so this tension needs careful consideration in protocol design. Making networks unmanageable to mitigate PM is not an acceptable outcome, but ignoring PM would go against the consensus documented here. An appropriate balance will emerge over time as real instances of this tension are considered.

Finally, the IETF, as a standards development organisation, does not control the implementation or deployment of our specifications (though IETF participants do develop many implementations), nor does the IETF standardise all layers of the protocol stack. Moreover, the non-technical (e.g., legal and political) aspects of mitigating pervasive monitoring are outside of the scope of the IETF. The broader Internet community will need to step forward to tackle PM, if it is to be fully addressed.

To summarise: current capabilities permit some actors to monitor content and metadata across the Internet at a scale never before seen. This pervasive monitoring is an attack on Internet privacy. The IETF will strive to produce specifications that mitigate pervasive monitoring attacks.

3. Process Note

In the past, architectural statements of this sort, e.g., [RFC1984] and [RFC2804], have been published as joint products of the Internet Engineering Steering Group (IESG) and the Internet Architecture Board (IAB). However, since those documents were published, the IETF and IAB have separated their publication "streams" as described in [RFC4844] and [RFC5741]. This document was initiated after discussions in both the IESG and IAB, but is published as an IETF-stream consensus document, in order to ensure that it properly reflects the consensus of the IETF community as a whole.

4. Security Considerations

This document is entirely about privacy. More information about the relationship between security and privacy threats can be found in [RFC6973]. Section 5.1.1 of [RFC6973] specifically addresses surveillance as a combined security-privacy threat.

5. Acknowledgements

We would like to thank the participants of the IETF 88 technical plenary for their feedback. Thanks in particular to the following for useful suggestions or comments: Jari Arkko, Fred Baker, Marc Blanchet, Tim Bray, Scott Brim, Randy Bush, Brian Carpenter, Benoit Claise, Alissa Cooper, Dave Crocker, Spencer Dawkins, Avri Doria, Wesley Eddy, Adrian Farrel, Joseph Lorenzo Hall, Phillip Hallam-Baker, Ted Hardie, Sam Hartmann, Paul Hoffman, Bjoern Hoehrmann, Russ Housley, Joel Jaeggli, Stephen Kent, Eliot Lear, Barry Leiba, Ted Lemon, Subramanian Moonesamy, Erik Nordmark, Pete Resnick, Peter Saint-Andre, Andrew Sullivan, Sean Turner, Nicholas Weaver, Stefan Winter, and Lloyd Wood. Additionally, we would like to thank all those who contributed suggestions on how to improve Internet security and privacy or who commented on this on various IETF mailing lists, such as the ietf@ietf.org and the perpass@ietf.org lists.

6. Informative References

[IETF88Plenary]
IETF, "IETF 88 Plenary Meeting Materials", November 2013, <http://www.ietf.org/proceedings/88/>.

[RFC1984] IAB, IESG, Carpenter, B., and F. Baker, "IAB and IESG Statement on Cryptographic Technology and the Internet", RFC 1984, August 1996.

[RFC2804] IAB and IESG, "IETF Policy on Wiretapping", RFC 2804, May 2000.

[RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, July 2003.

[RFC4844] Daigle, L. and Internet Architecture Board, "The RFC Series and RFC Editor", RFC 4844, July 2007.

[RFC4949] Shirey, R., "Internet Security Glossary, Version 2", RFC 4949, August 2007.

[RFC5741] Daigle, L., Kolkman, O., and IAB, "RFC Streams, Headers, and Boilerplates", RFC 5741, December 2009.

[RFC6962] Laurie, B., Langley, A., and E. Kasper, "Certificate Transparency", RFC 6962, June 2013.

[RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", RFC 6973, July 2013.

Authors' Addresses

Stephen Farrell
Trinity College Dublin
Dublin 2
Ireland

Phone: +353-1-896-2354
EMail: stephen.farrell@cs.tcd.ie

Hannes Tschofenig
ARM Ltd.
6060 Hall in Tirol
Austria

EMail: Hannes.tschofenig@gmx.net
URI: http://www.tschofenig.priv.at